REMARKS

I. Introduction

Claims 1-21 are all the claims pending in the application, and claims 1-21 have been examined. Claims 1-21 stand rejected under 35 U.S.C. § 103(a) as allegedly being unpatentable over the Chen et al. publication, *Performance Comparison of Three Alternatives of Distributed Multidatabase Systems: A Global Query Perspective* (hereinafter "Chen"), in view Jindal et al., U.S. Patent No. 6,324,580 (hereinafter "Jindal").

Applicants overcome the rejection of claims 1-21 as follows.

II. Claim Rejections -- 35 U.S.C. § 103(a)

A. Claims 1, 7, and 13

The Examiner acknowledges that Chen fails to teach or suggest "selecting a server to process the request . . . based on whether the server can satisfy the request for data", as recited in claims 1, 7, and 13 (*see* Office Action, page 3). Jindal fails to make up for this acknowledged deficiency of Chen.

The Examiner alleges that Jindal teaches this feature because Jindal describes a selected policy for choosing a preferred server such as a least-loaded server or the closest server (col. 6, lines 32-46). However, neither a determination of the least-loaded server nor of the closest server ensures that a server can satisfy the request for data, as recited in claim 1 (*see also* claims 7 and 13). For a system that can support a plurality of heterogeneous data sources (*see* claims 1, 7, and 13), it is necessary to determine that a server of the system can satisfy a request for data. *See, e.g.,* Applicants' Fig. 7. For example, if a server is not of the proper type, it won't be able to

3

satisfy the request for data, regardless of its load. *See, e.g.,* Applicants' page 45, line 21 to page

46, line 1. Furthermore, even determining whether a server is currently operational (*see, e.g.,*

col. 6, lines 47-55 of Jindal) does not ensure that a server <u>can satisfy the request for data</u> because

an operational server may still be unable to satisfy the request for data.

Jindal describes load balancing requests for a replicated service or application among a

plurality of servers operating instances of the replicated service or application. *See, e.g.,*

Abstract of Jindal. Thus, aside from their current load, the servers in Jindal are the same. *Id.*

Because the servers in Jindal all offer the same (replicated) service, there is no need to select a

sever "based on whether the server can satisfy the request for data". Indeed, in Jindal, if one

server can satisfy the request for data then all the other servers can as well.

Furthermore, Applicants note that to establish a *prima facie* case of obviousness, there

must be some suggestion or motivation, either in the references themselves or in the knowledge

generally available to one of ordinary skill in the art, to modify the reference or to combine

reference teachings. *See* MPEP § 2143.

Furthermore, as the Federal Circuit just recently reminded us, the USPTO is held to a

rigorous standard when trying to show that an invention would have been obvious in view of the

combination of two or more references. *See, In re Sang Su Lee*, 2002 U.S. App. LEXIS 855, *10

(Fed. Cir. 2002), *citing, e.g., In re Dembiczak*, 175 F.3d 994, 999, 50 USPQ2d 1614, 1617 (Fed.

Cir. 1999) ("Our case law makes clear that the best defense against the subtle but powerful

attraction of a hindsight-based obviousness analysis is rigorous application of the requirement for

a showing of the teaching or motivation to combine prior art references."). The Federal Circuit

goes on to emphasize that the "need for specificity pervades this authority." *In re Sang Su Lee* at *10-*11 (emphasis added) (*citing In re Kotzab*, 217 F.3d 1365, 1371, 55 USPQ2d 1313, 1317 (Fed. Cir. 2000) ("particular findings must be made as to the reason the skilled artisan, with no knowledge of the claimed invention, would have selected these components for combination in the manner claimed")).

Applicants respectfully submit that the current grounds of rejection do not satisfy the aforementioned standard for demonstrating that the claimed invention would have been obvious in view of the combination of Chen in view of Jindal.

Specifically, the Examiner asserts that "it would have been obvious to one of ordinary skill in the art at the time the invention was made to modify Chen by including selecting a server to process the request based on a load of the server and based on whether the server can satisfy the request for data, as taught by Jindal, so the computer server capable of efficiently satisfying the needs of a limited number of clients". However, there does not appear to be any suggestion or motivation, either in the references themselves or in the knowledge generally available to one of ordinary skill in the art, to modify the reference or to combine reference teachings. For example, Chen presents a comparative study of three common relational multidatabase approaches from the perspective of global query performance. However, satisfying the needs of a limited number of clients is not an issue taught or suggested by Chen. Indeed, the issue only arises in Jindal where a plurality of servers provide instances of the same service, but where a server may only be able to satisfy the needs of a limited number of clients.

Furthermore, Jindal describes load balancing in general terms relating to replicated services. Jindal is silent with respect to searching for data in heterogeneous data sources, as recited in claims 1-21. Indeed, to the extend that the servers of Jindal "mirror" one another, they are certainly not heterogeneous.

Thus, absent impermissible hindsight, the Examiner has provided no indication as to why one of ordinary skill in the art would have been motivated from the disclosure of Jindal, relating to balancing client requests for a service replicated among a plurality of servers , to modify the general multidatabase approach described in Chen in order to obtain Applicants' claimed invention.

Additionally, Applicants amend claims 1, 7, and 13 to further clarify that the selected server is connected to one or more heterogeneous datastores. *See, e.g.,* Applicants' Fig. 7. Jindal describes a plurality of servers that each offer the same (*i.e.,* homogeneous) service. *See, e.g.,* claim 1; Fig. 1 of Jindal. Jindal fails to teach or suggest a server connected to one or more heterogeneous datastores. Thus, Jindal fails to teach or suggest "selecting a server to process the request based on a load of the server and based on whether the server can satisfy the request for data, said server connected to one or more heterogeneous datastores", as recited in claim 1 (*see also* claims 7 and 13). Chen merely describes the general concept of a federated multidatabase system, and therefore, Chen fails to make up for the deficiencies of Jindal.

B. Claims 2-6, 8-12, and 14-21

Claims 2-6, 8-12, and 14-21 are patentable over the combination of Chen and Jindal at least by virtue of their respective dependency from independent claims 1, 7, and 13, as well as the additional features recited therein.

For example, claims 3, 9, and 15 recite "forwarding additional requests for similar data to the selected server". The Examiner asserts that Jindal discloses this recited feature at col. 6, line 4 to col. 7, line 21. Jindal describes that a policy for selecting a preferred server may be in effect for a short time (*e.g.,* 60 seconds) or a long time (e.g., 1 day). *See, e.g.,* col. 7, lines 11-21 of Jindal. However, this duration of the preferred server described in Jindal does not teach or suggest "forwarding additional requests for similar data to the selected server". Whether a request is <u>for similar data</u> is not a consideration taught or suggested by Jindal. Indeed, since Jindal relates to a replicated service or application, all requests (for the service or application) would be for similar data.

Furthermore, claims 5, 11, and 17 recite "upon receiving a request to add another server, connecting the server to an existing server in the server hierarchy based on a number of connections of the existing server. The Examiner asserts that Chen discloses these claimed features at page 54/col. 2, line 25 to page 55/col. 1, line 5 and Fig. 3. The portion of Chen relied on upon by the Examiner discusses a middleware (software) approach to conducting global queries in a multidatabase system. However, the middleware discussion of Chen fails to teach or suggest "receiving a request to add another server" or "connecting the server to an existing server in the server hierarchy based on a number of connections of the existing server". Indeed, Fig. 3 of Chen fails to even illustrate (or teach or suggest) a "server hierarchy" in that the

7

middleware/software is not a server. Thus, it follows that Chen (contrary to the Examiner's assertion), also fails to teach or suggest "upon receiving a request to delete an existing server in the hierarchy, deleting that server", as recited in claims 6, 12, and 18.
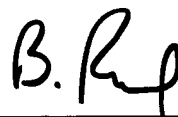
For at least the above exemplary reasons, claims 1-21 are patentable over the combination of Chen and Jindal.

III. Conclusion

In view of the above, reconsideration and allowance of this application are now believed to be in order, and such actions are hereby solicited. If any points remain in issue which the Examiner feels may be best resolved through a personal or telephone interview, the Examiner is kindly requested to contact the undersigned attorney at the telephone number listed below.

The USPTO is directed and authorized to charge all required fees, except for the Issue Fee and the Publication Fee, to Deposit Account No. 19-4880. Please also credit any overpayments to said Deposit Account.

Respectfully submitted,

B. R—

SUGHRUE MION, PLLC   Billy Carter Raulerson
Telephone: (202) 293-7060   Registration No. 52,156
Facsimile: (202) 293-7860

WASHINGTON OFFICE

23373

PATENT TRADEMARK OFFICE

Date: February 28, 2003
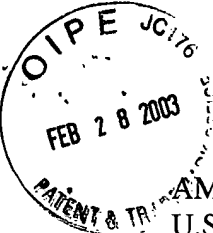
APPENDIX

VERSION WITH MARKINGS TO SHOW CHANGES MADE

IN THE CLAIMS:

The claims are amended as follows:

1.   (Amended) A method for searching for data in one or more heterogeneous data sources within a computer system, the method comprising the steps of:

receiving a request for data at a federated data source; and

selecting a server to process the request based on a load of the server and based on whether the server can satisfy the request for data, said server connected to one or more heterogeneous datastores.

7.   (Amended) An apparatus for searching for data in one or more heterogeneous data sources, comprising:

a computer system having one or more heterogeneous data sources; and

one or more computer programs, performed by the computer system, for receiving a request for data at a federated data source and selecting a server to process the request based on a load of the server and based on whether the server can satisfy the request for data, said server connected to one or more heterogeneous datastores.

13.   (Amended) An article of manufacture comprising a program storage medium readable by a computer system and embodying one or more instructions executable by the

computer system to perform method steps for searching for data in one or more heterogeneous

data sources within a computer system, the method comprising the steps of:

receiving a request for data at a federated data source; and

selecting a server to process the request based on a load of the server and based on

whether the server can satisfy the request for data, said server connected to one or more

heterogeneous datastores.